JMLR: Workshop and Conference Proceedings 1–25

arXiv:1605.08722v1 [cs.LG] 27 May 2016

An algorithm with nearly optimal pseudo-regret for both stochastic and adversarial bandits

Peter Auer * AUER@UNILEOBEN.AC.AT
Chair for Information Technology
Montanuniversitaet Leoben
Franz-Josef-Strasse 18, A-8700 Leoben, Austria

Chao-Kai Chiang CHAOKAI@GMAIL.COM
University of California, Los Angeles
Computer Science Department
4732 Boelter Hall
Los Angeles, CA 90095

Abstract

We present an algorithm that achieves almost optimal pseudo-regret bounds against adversarial and stochastic bandits. Against adversarial bandits the pseudo-regret is $O\left(\sqrt{Kn\log n}\right)$ and against stochastic bandits the regret is $O\left(\sum_i (\log n)/\Delta_i\right)$. We also show that no algorithm with $O\left(\log n\right)$ pseudo-regret against stochastic bandits can achieve $\tilde{O}\left(\sqrt{n}\right)$ expected regret against adaptive adversarial bandits. This complements previous results of Bubeck and Slivkins (2012) that show $\tilde{O}\left(\sqrt{n}\right)$ expected adversarial regret with $O\left((\log n)^2\right)$ stochastic pseudo-regret.

1. Introduction

We consider the multi-armed bandit problem, which is the most basic example of a sequential decision problem with an exploration-exploitation trade-off. In each time step $t = 1, 2, \ldots, n$, the player has to play an arm $I_t \in \{1, \ldots, K\}$ from this fixed finite set and receives reward $x_{I_t}(t) \in [0, 1]$ depending on its choice[1]. The player observes only the reward of the chosen arm, but not the rewards of the other arms $x_i(t)$, $i \neq I_t$. The player's goal is to maximize its total reward $\sum_{t=1}^n x_{I_t}(t)$, and this total reward is compared to the best total reward of a single arm, $\sum_{t=1}^n x_i(t)$. To identify the best arm the player needs to explore all arms by playing them, but it also needs to limit this exploration to often play the best arm. The optimal amount of exploration constitutes the exploration-exploitation trade-off.

Different assumptions on how the rewards $x_i(t)$ are generated have led to different approaches and algorithms for the multi-armed bandit problem. In the original formulation (Robbins, 1952) it is assumed that the rewards are generated independently at random, governed by fixed but unknown probability distributions with means μ_i for each arm $i = 1, \ldots, K$. This type of bandit problem is called *stochastic*. The other type of bandit problem that we consider in this paper is called non-stochastic or *adversarial* (Auer et al., 2002b). Here the rewards may be selected arbitrarily by

* Accepted for presentation at the Conference on Learning Theory (COLT) 2016.

1. We assume that the player knows the total number of time steps n.

© P. Auer & C.-K. Chiang.

an adversary and the player should still perform well for any selection of rewards. An extensive overview of multi-armed bandit problems is given in (Bubeck and Cesa-Bianchi, 2012).

A central notion for the analysis of stochastic and adversarial bandit problems is the regret $R(n)$, the difference between the total reward of the best arm and the total reward of the player:

$$R(n) = \max_{1 \le i \le K} \sum_{t=1}^{n} x_i(t) - \sum_{t=1}^{n} x_{I_t}(t).$$

Since the player does not know the best arm beforehand and needs to do exploration, we expect that the total reward of the player is less than the total reward of the best arm. Thus the regret is a measure for the cost of not knowing the best arm. In the analysis of bandit problems we are interested in high probability bounds on the regret or in bounds on the expected regret. Often it is more convenient, though, to analyze the pseudo-regret

$$\overline{R}(n) = \max_{1 \le i \le K} \mathbb{E}\left[\sum_{t=1}^{n} x_i(t) - \sum_{t=1}^{n} x_{I_t}(t)\right]$$

instead of the expected regret

$$\mathbb{E}\left[R(n)\right] = \mathbb{E}\left[\max_{1 \le i \le K} \sum_{t=1}^{n} x_i(t) - \sum_{t=1}^{n} x_{I_t}(t)\right].$$

While the notion of pseudo-regret is weaker than the expected regret with $\overline{R}(n) \le \mathbb{E}\left[R(n)\right]$, bounds on the pseudo-regret imply bounds on the expected regret for adversarial bandit problems with *oblivious* rewards $x_i(t)$ selected independently from the player's choices. The pseudo-regret also allows for refined bounds in stochastic bandit problems.

1.1. Previous results

For adversarial bandit problems, algorithms with high probability bounds on the regret are known (Bubeck and Cesa-Bianchi, 2012, Theorem 3.3): with probability $1 - \delta$,

$$R_{\mathrm{adv}}(n) = O\left(\sqrt{n \log(1/\delta)}\right).$$

For stochastic bandit problems, several algorithms achieve logarithmic bounds on the pseudo-regret, e.g. Auer et al. (2002a):

$$\overline{R}_{\mathrm{sto}}(n) = O\left(\log n\right).$$

Both of these bounds are known to be best possible.

While the result for adversarial bandits is a worst-case — and thus possibly pessimistic — bound that holds for any sequence of rewards, the strong assumptions for stochastic bandits may sometimes be unjustified. Therefore an algorithm that can adapt to the actual difficulty of the problem is of great interest. The first such result was obtained by Bubeck and Slivkins (2012), who developed the SAO algorithm that with probability $1 - \delta$ achieves

$$R_{\mathrm{adv}}(n) \le O\left((\log n)\sqrt{n \log(n/\delta)}\right)$$

regret for adversarial bandits and

$$\overline{R}_{\mathrm{sto}}(n) = O\left((\log n)^2\right)$$

pseudo-regret for stochastic bandits.

It has remained as an open question if a stochastic pseudo-regret of order $O\left((\log n)^2\right)$ is necessary or if the optimal $O(\log n)$ pseudo-regret can be achieved while maintaining an adversarial regret of order $\sqrt{n}$.

1.2. Summary of new results

We give a twofold answer to this open question. We show that stochastic pseudo-regret of order $O\left((\log n)^2\right)$ is necessary for a player to achieve high probability adversarial regret of order $\sqrt{n}$ against an oblivious adversary, and to even achieve expected regret of order $\sqrt{n}$ against an adaptive adversary. But we also show that a player can achieve $O(\log n)$ stochastic pseudo-regret and $\tilde{O}(\sqrt{n})$ adversarial *pseudo-regret* at the same time. This gives, together with the results of (Bubeck and Slivkins, 2012), a quite complete characterization of algorithms that perform well both for stochastic and adversarial bandit problems.

More precisely, for any player with stochastic pseudo-regret bound of order $O\left((\log n)^\beta\right)$, $\beta < 2$, and any $\epsilon > 0$, $\alpha < 1$, there is an adversarial bandit problem for which the player suffers $\Omega(n^\alpha)$ regret with probability $\Omega(n^{-\epsilon})$. Furthermore, there is an adaptive adversary against which the player suffers $\Omega(n^\alpha)$ expected regret. Secondly, we construct an algorithm with

$$\overline{R}_{\mathrm{sto}}(n) = O(\log n)$$

and

$$\overline{R}_{\mathrm{adv}}(n) = O\left(\sqrt{n \log n}\right).$$

At first glance these two results may appear contradictory for $\alpha - \epsilon > 1/2$, as the lower bound seems to suggest a pseudo-regret of $\Omega(n^{\alpha-\epsilon})$. This is not the case, though, since the regret may also be negative. Indeed, consider an adversarial multi-armed bandit that initially gives higher rewards for one arm, and from some time step on gives higher rewards for a second arm. A player that detects this change and initially plays the first arm and later the second arm, may outperform both arms and achieve negative regret. But if the player misses the change and keeps playing the first arm, it may suffer large regret against the second arm.

In our analysis we use both mechanisms. For the lower bound on the pseudo-regret we show that a player with little exploration (which is necessary for small stochastic pseudo-regret) will miss such a change with significant probability and then will suffer large regret. For the upper bound we explicitly compensate possible large regret that occurs with small probability by negative regret that occurs with sufficiently large probability. For the lower bound on the expected regret we construct an adaptive adversary that prevents such negative regret. Consequently, our results exhibit one of the rare cases where there is a significant gap between the achievable pseudo-regret and the achievable expected regret.

The explicit consideration of negative regret is one of the technical contributions of this work. Another, maybe even more significant contribution, is a weak testing scheme for non-stochastic arms. This weak testing scheme is necessary since $O(\log n)$ stochastic pseudo-regret allows only for very little exploration. Each individual weak test has a constant false positive rate (predicting

a non-stochastic arm although the arm is stochastic) and a constant false negative rate (missing a non-stochastic arm). To avoid classifying a stochastic arm as non-stochastic, an arm is classified as non-stochastic only after $O(\log n)$ positive tests. This reduces the false positive rate of a decision to acceptable $O(1/n)$. Conversely, this delayed detection needs to be accounted for in the regret analysis when the arms are indeed non-stochastic.

2. Definitions and statement of results

In a multi-armed bandit problem with arms $i = 1, \ldots, K$ the interaction of a player with its environment is governed by the following protocol:

For time steps $t = 1, \ldots, n$:

1. The player chooses an arm $I_t \in \{1, \ldots, K\}$, possibly using randomization.
2. The player receives and observes the reward $x_{I_t}(t)$.
   It does not observe the reward from any other arm $i \neq I_t$.

The player's choice I_t may depend only on information available at this time, namely $I_1, \ldots, I_{t-1}$ and $x_{I_1}(1), \ldots, x_{I_{t-1}}(t-1)$. If the bandit problem is stochastic, then the rewards $x_i(t)$ are generated independently at random. If the bandit problem is adversarial, then the rewards are generated arbitrarily by an adversary. We assume that all rewards $x_i(t) \in [0, 1]$ and that the number of time steps n is known to the player.

2.1. Stochastic multi-armed bandit problems

In a stochastic multi-armed bandit problem the rewards for each arm i are generated by a fixed but unknown probability distribution ν_i on $[0, 1]$. All rewards $x_i(t)$, $1 \leq i \leq K$, $1 \leq t \leq n$, are generated independently at random with $x_i(t) \sim \nu_i$.

Important quantities are the average rewards of the arms, $\mu_i = \mathbb{E}[x_i(t)]$, the average reward of the best arm $\mu^* = \max_i \mu_i$, and the resulting gaps $\Delta_i = \mu^* - \mu_i$.

The goal of the player is to achieve low pseudo-regret which for a stochastic bandit problem can be written as

$$\overline{R}_{\mathrm{sto}}(n) = \max_{1 \leq i \leq K} \mathbb{E}\left[\sum_{t=1}^{n} x_i(t) - \sum_{t=1}^{n} x_{I_t}(t)\right] = \sum_{i=1}^{K} \Delta_i \mathbb{E}\left[T_i(n)\right],$$

where $T_i(n) = \#\{1 \leq t \leq n : I_t = i\}$ is the number of plays of arm i. It can be shown (Auer et al., 2002a) that — among others — upper confidence bound algorithms achieve

$$\mathbb{E}\left[T_i(n)\right] = O\left(\frac{\log n}{\Delta_i^2}\right)$$

for any arm i with $\Delta_i > 0$ such that

$$\overline{R}_{\mathrm{sto}}(n) = O\left(\sum_{i:\Delta_i > 0} \frac{\log n}{\Delta_i}\right).$$

It can be even shown that for arms i with $\Delta_i > 0$,

$$T_i(n) = O\left(\frac{\log(n/\delta)}{\Delta_i^2}\right)$$

with probability $1-\delta$ when n is known to the player.

2.2. Adversarial multi-armed bandit problems

In adversarial bandit problems the rewards are selected by an adversary. If this is done beforehand (before the player interacts with the environment), then the adversary is called *oblivious* as the selection of rewards is independent from the arms I_t chosen by the player. In this case any upper bound on the pseudo-regret that holds for any selection of rewards is also an upper bound on the expected regret.

If the selection of rewards $x_i(t)$, $1 \le i \le K$, depends on which arms $I_1, \dots, I_{t-1}$ the player has chosen in the past, then the adversary is called *adaptive*. In this case a bound on the pseudo-regret does not necessarily translate into a bound on the expected regret. Nevertheless, strong bounds on the regret against an adaptive adversary are known for the EXP3.P algorithm (Auer et al., 2002b):

Theorem 1 (Bubeck and Cesa-Bianchi, 2012, Theorem 3.3) *When* EXP3.P *is run with appropriate parameters depending on n, K, and δ, then with probability $1-\delta$ its regret satisfies*

$$R_{\text{ada}}(n) = O\left(\sqrt{nK\log(K/\delta)}\right).$$

2.3. Results

First, we state our lower bounds for oblivious and adaptive adversaries.

Theorem 2 *Let $\alpha < 1$, $\epsilon > 0$, $\beta < 2$, and $C > 0$. Consider a player that achieves pseudo-regret*

$$\overline{R}_{\text{sto}}(n) \le C(\log n)^{\beta}$$

for any stochastic bandit problem with two arms and gap $\Delta = 1/8$. Then for large enough n there is an adversarial bandit problem with two arms and an oblivious adversary such that the player suffers regret

$$R_{\text{obl}}(n) \ge n^{\alpha}/8 - 4\sqrt{n\log n}$$

with probability at least $1/(16n^{\epsilon}) - 2/n^2$. Furthermore, there is an adversarial bandit problem with two arms and an adaptive adversary such that the player suffers expected regret

$$\mathbb{E}\left[R_{\text{ada}}(n)\right] \ge \frac{n^{\alpha-\epsilon}}{128} - 3\sqrt{n\log n}.$$

In Section 3 we present our SAPO algorithm (Stochastic and Adversarial Pseudo-Optimal) that achieves optimal pseudo-regret in stochastic bandit problems and nearly optimal pseudo-regret in adversarial bandit problems. Its performance is summarized in the following theorem.

Theorem 3 *For large enough n and any $\delta > 0$, algorithm SAPO achieves the following bounds for suitable constants C_{sto}, C_{adv}, and C_{1b}:*

- *For stochastic bandit problems with gaps Δ_i such that $C_{1\mathrm{b}} \sum_{i:\Delta_i>0} \frac{\log(n/\delta)}{\Delta_i} \le \sqrt{nK\log(n/\delta)}$,*

$$T_i(n) \le C_{\mathrm{sto}} \frac{\log(n/\delta)}{\Delta_i^2}$$

with probability $1-\delta$ for any arm i with $\Delta_i > 0$, and thus

$$\overline{R}_{\mathrm{sto}}(n) \le C_{\mathrm{sto}} \sum_{i:\Delta_i>0} \frac{\log(n/\delta)}{\Delta_i} + \delta n.$$

- *For adversarial bandit problems*

$$\overline{R}_{\mathrm{ada}}(n) \le C_{\mathrm{adv}} K \sqrt{n\log(n/\delta)} + \delta n.$$

Remark 4 *Our bound for adversarial bandit problems shows a worse dependency on K than Theorem 1. This is an artifact of our current analysis and can be improved to a bound $\overline{R}_{\mathrm{ada}}(n) = O\left(\sqrt{nK\log(n/\delta)}\right)$.*

2.4. Comparison with related work

Bubeck and Slivkins (2012) show for their SAO algorithm that with probability $1-\delta$,

$$\sum_{i=1}^{K} \Delta_i T_i(n) \le O\left(\frac{K\log K(\log n/\delta)^2}{\Delta}\right)$$

for stochastic bandits where $\Delta = \min_{i:\Delta_i>0} \Delta_i$, and

$$R_{\mathrm{ada}}(n) \le O\left((\log K)(\log n)\sqrt{nK\log n/\delta}\right)$$

for adaptive adversarial bandits. While our bounds in Theorem 3 are somewhat tighter, in particular showing the optimal dependency on the gaps Δ_i for stochastic bandits, we have only a result on the pseudo-regret for adversarial bandits. We conjecture though, that our analysis can be used to construct an algorithm that with probability $1-\delta$ achieves $T_i(n) \le O\left((\log n/\delta)^2/\Delta_i^2\right)$ for stochastic bandits and $R_{\mathrm{ada}}(n) \le O\left((\log K)(\log n)\sqrt{nK\log n/\delta}\right)$ for adaptive adversarial bandits.

Our SAPO algorithm follows the general strategy of the SAO algorithm by essentially employing an algorithm for stochastic bandit problems that is equipped with additional tests to detect non-stochastic arms. A different approach is taken in (Seldin and Slivkins, 2014): here the starting point is an algorithm for adversarial bandit problems that is modified by adding an additional exploration parameter to achieve also low pseudo-regret in stochastic bandit problems. While this approach has not yet allowed for the tight $O(\log n)$ regret bound in stochastic bandit problems (they achieve a $O\left(\log^3 n\right)$ bound), the approach is quite flexible and more generally applicable than the SAO and SAPO algorithms.

2.5. Proof sketch of the lower bound (Theorem 2)

We present here the main idea of the proof. The proof itself is given in Appendix B.

We consider a stochastic bandit problem with constant reward $x_1(t) = 1/2$ for arm 1 and Bernoulli rewards with $\mu_2 = 1/2 - \Delta$ for arm 2, $\Delta = 1/8$. We divide the time steps into phases of increasing length $L_j = 3^j n^\alpha$, $j = 0, \dots, J$ with $J = \Omega(\log n)$. Since the pseudo-regret of the player is $O\left((\log n)^\beta\right)$, there is a phase j^* where the expected number of plays of arm 2 in this phase is $O\left((\log n)^{\beta-1}\right)$.

We construct an oblivious adversarial bandit by modifying the Bernoulli distribution of arm 2 in phase j^* and beyond by setting $\mu_2 = 1/2 + \Delta$. By this modification arm 2 gives larger total reward than arm 1.

Because of the limited number of plays in phase j^*, a standard argument shows that the player will not detect this modification during phase j^* with probability $\exp\{-O(\log^{\beta-1} n)\} = \Omega(n^{-\epsilon})$. When the modification is not detected during phase j^*, then in this phase the player suffers roughly regret ΔL_{j^*} against arm 2. This is not compensated by negative regret against arm 2 in previous phases since $\Delta \sum_{j=0}^{j^*-1} L_j \le \Delta L_{j^*}/2$. Thus in this case the overall regret of the player against arm 2 is roughly $\Delta L_{j^*}/2 = \Omega(n^\alpha)$.

In a very similar way we can construct also an adaptive adversarial bandit: As for the oblivious bandit, we set $\mu_2 = 1/2 + \Delta$ in phase j^*. If the player chooses arm 2 only $C(\log n)^{\beta-1}$ times in phase j^*, then we keep $\mu_2 = 1/2 + \Delta$ also for the remaining phases. As for the oblivious bandit this happens with probability $\Omega(n^{-\epsilon})$ and gives regret $\Omega(n^\alpha)$. To avoid negative regret, we switch back to $\mu_2 = 1/2 - \Delta$, as soon as there more than $C(\log n)^{\beta-1}$ plays of arm 2 in phase j^*. In this case the reward of the algorithm is roughly $n/2 + C\Delta(\log n)^{\beta-1}$ such that in this case $R(n) \ge -C\Delta(\log n)^{\beta-1}$. Hence the expected regret is $\mathbb{E}[R(n)] \ge \Omega(n^{\alpha-\epsilon}) - C\Delta(\log n)^{\beta-1} = \Omega(n^{\alpha-\epsilon})$.

3. The SAPO algorithm

In its core the algorithm is an elimination procedure for stochastic bandits that is augmented by tests safeguarding against non-stochastic arms. If there is sufficient evidence for non-stochastic arms, then the algorithm switches to the adversarial bandit algorithm EXP3.P, starting with the current time step.

The algorithm maintains a set of active arms $\mathcal{A}$ and a set of supposedly suboptimal "bad" arms $\mathcal{B}$. For each arm i it maintains the sample mean $\hat{\mu}_i(s)$,

$$\hat{\mu}_i(s) = \frac{1}{T_i(s)} \sum_{t=1}^{s} x_i(t) \mathbb{I}[I_t = i],$$
$$T_i(s) = \sum_{t=1}^{s} \mathbb{I}[I_t = i],$$

and also an unbiased estimate to deal with non-stochastic arms,

$$\bar{\mu}_i(s) = \frac{1}{s} \sum_{t=1}^{s} x_i(t) \frac{\mathbb{I}[I_t = i]}{p_i(t)},$$

Algorithm 1 : SAPO

Input: Number of arms K, number of rounds $n \geq K$, and confidence parameter δ.

Initialization: All arms are active, $\mathcal{A}(0) = \{1, \ldots, K\}$, $\mathcal{B}(0) = \emptyset$.

For $t = 1, \ldots, n$:

1. (a) If there is an arm $i \in \mathcal{A}(t-1)$ with $\bar{\mu}_i(t-1) \notin [\overline{\mathrm{lcb}}_i(t-1), \overline{\mathrm{ucb}}_i(t-1)]$, then switch to EXP3.P.
   (b) If $\sum_{s=1}^{t-1} [\mathrm{lcb}^*(s) - x_{I_s}(s)] > C_{1\mathrm{b}} \sqrt{Kn \log(n/\delta)}$, then switch to Exp3.P.
2. Evict arms from $\mathcal{A}$:
   (a) Let $B(t) = \{i \in \mathcal{A}(t-1) : \; T_i(t-1) \geq C_{\mathrm{init}} \cdot \log(n/\delta) \wedge \hat{\mu}_i(t-1) + C_{\mathrm{gap}} \cdot \mathrm{width}_i(t-1) < \mathrm{lcb}^*(t-1)\}$, $\mathcal{A}(t) = \mathcal{A}(t-1) \setminus B(t)$, $\mathcal{B}(t) = \mathcal{B}(t-1) \cup B(t)$.
   (b) For all $i \in B(t)$ set $\tilde{\mu}_i = \hat{\mu}_i(t-1)$, $\tilde{\Delta}_i = C_{\mathrm{gap}} \cdot \mathrm{width}_i(t-1)$, $n_i(t) = t$, $L_i(t) = L_i^0 := \lceil C_{\mathrm{p}} K / \tilde{\Delta}_i^2 \rceil$, and $E_i(t) = 0$.
3. Choose $I_t = i$ with probabilities

$$p_i(t) = \begin{cases} L_i^0 / (K L_i(t)) & \text{for } i \in \mathcal{B}(t) \\ \left(1 - \sum_{j \in \mathcal{B}(t)} p_j(t)\right) \Big/ |\mathcal{A}(t)| & \text{for } i \in \mathcal{A}(t) \end{cases}$$

4. Test and update all arms $i \in \mathcal{B}(t)$:
   (a) If $\exists s : n_i(t) \leq s \leq t : \hat{D}_i(s,t) \geq C_{4\mathrm{a}} \tilde{\Delta}_i L_i(t) p_i(t)$,
   (b) then $n_i(t+1) = t+1$, $L_i(t+1) = \max\{L_i(t)/2, L_i^0\}$, and $E_i(t+1) = E_i(t) + 1$,
   (c) if $E_i(t+1) = E^0 := \lceil C_{\mathrm{E}} \cdot \log(n/\delta) \rceil$, then switch to EXP3.P;
   (d) else if $t = n_i(t) + L_i(t) - 1$ then $n_i(t+1) = t+1$, $L_i(t+1) = 2L_i(t)$, and $E_i(t+1) = E_i(t)$;
   (e) else $n_i(t+1) = n_i(t)$, $L_i(t+1) = L_i(t)$, and $E_i(t+1) = E_i(t)$.

where $p_i(t)$ is the probability of choosing arm i at time t. Confidence bounds[2] around the estimated means are used to evict arms from the active set $\mathcal{A}$,

$$\begin{aligned} \mathrm{lcb}_i(s) &= \max\{\mathrm{lcb}_i(s-1), \hat{\mu}_i(s) - \mathrm{width}_i(s)\}, \\ \overline{\mathrm{lcb}}_i(s) &= \max\{\overline{\mathrm{lcb}}_i(s-1), \bar{\mu}_i(s) - \overline{\mathrm{width}}(s)\}, \\ \overline{\mathrm{ucb}}_i(s) &= \min\{\overline{\mathrm{ucb}}_i(s-1), \bar{\mu}_i(s) + \overline{\mathrm{width}}(s)\}, \\ \mathrm{lcb}^*(s) &= \max_{1 \leq i \leq K} \max\{\mathrm{lcb}_i(s), \overline{\mathrm{lcb}}_i(s)\}, \\ \mathrm{width}_i(s) &= \sqrt{C_{\mathrm{w}} \log(n/\delta) / T_i(s)}, \\ \overline{\mathrm{width}}(s) &= \sqrt{C_{\mathrm{w}} K \log(n/\delta) / s}. \end{aligned}$$

2. We start with $\mathrm{lcb}_i(0) = \overline{\mathrm{lcb}}_i(0) = 0$ and $\overline{\mathrm{ucb}}_i(0) = 1$.

Note that $\mathrm{lcb}_i(s)$, $\overline{\mathrm{lcb}}_i(s)$, and $\mathrm{lcb}^*(s)$ are non-decreasing and $\overline{\mathrm{ucb}}_i(s)$ are non-increasing. This reflects the intuition that confidence intervals should be shrinking and is used to safeguard against non-stochastic arms.

An arm i is evicted from $\mathcal{A}$ in **Step 2.a**, if it has a sufficient number of plays ($C_{\mathrm{init}} \cdot \log(n/\delta)$) for reasonably accurate estimates, and if its sample mean $\hat{\mu}_i(t-1)$ is significantly smaller than the optimal lower confidence bound $\mathrm{lcb}^*(t-1)$. The additional distance $C_{\mathrm{gap}} \cdot \mathrm{width}_i(t-1)$ is used to estimate the gap Δ_i. For evicted arms, in **Step 2.b** an estimate for the gap $\tilde{\Delta}_i$ and the current estimated mean are frozen, $\tilde{\mu}_i = \hat{\mu}_i(t-1)$. For stochastic bandits the accuracy of this estimate is proportional to the estimated gap $\tilde{\Delta}_i$. These quantities are used in the tests for detecting non-stochastic arms. Also the starting time $n_i(t)$ and the length $L_i(t) = L_i^0$ of the first testing phase (see below), as well as the number of detections $E_i(t) = 0$ are set.

Since SAPO needs to perform well also against adversaries, all choices of arms are randomized. In **Step 3** an active arm is chosen uniformly at random, or with some smaller probabilitya bad arm i is chosen where the probability depends on the length of its current testing phase $L_i(t)$. Choosing also bad arms is necessary to detect non-stochastic arms among the bad arms.

3.1. Tests for detecting non-stochastic arms

The most important test is in Step 4.a for detecting that a bad arm receives larger rewards than it should if it were stochastic. Such an arm could be optimal if the bandit problem is adversarial. The best way to view this test is by dividing the time steps of an evicted arm i into testing phases

$$\tau_{i,1}, \ldots, \tau_{i,2} - 1; \tau_{i,2}, \ldots, \tau_{i,3} - 1; \tau_{i,3}, \ldots, \tau_{i,4} - 1; \ldots$$

The first phase starts when arm i is evicted from $\mathcal{A}$. A phase k ends at time $\tau_{i,k+1} - 1$ if either the phase has exhausted its length (Step 4.d), or when the test in Step 4.a reports a detection.[3] Thus the length parameter $L_i(t)$ is only the maximal length of a phase and the phase may end earlier. In the notation of the algorithm $n_i(t)$ denotes the start of the current phase. Within a phase the probability $p_i(t)$ for choosing arm i is constant since the length parameter $L_i(t)$ does not change (Step 4.e). For notational convenience we denote by p_{ik} the probability for choosing arm i in its k-th testing phase, and by L_{ik} the corresponding length parameter,

$$\begin{aligned} p_i(t) &= p_{ik} \text{ for } i \in \mathcal{B}(t) \text{ and } \tau_{i,k} \le t < \tau_{i,k+1}, \\ L_i(t) &= L_{ik} \text{ for } i \in \mathcal{B}(t) \text{ and } \tau_{i,k} \le t < \tau_{i,k+1}, \\ n_i(t) &= \tau_{i,k} \text{ for } i \in \mathcal{B}(t) \text{ and } \tau_{i,k} \le t < \tau_{i,k+1}. \end{aligned}$$

Now the test in **Step 4.a** checks if a bad arm i has received significantly more rewards in the current phase then expected, given the estimated mean $\tilde{\mu}_i$, the maximal phase length $L_i(t)$ and the probability for choosing arm i, $p_i(t)$, where

$$\hat{D}_i(s_1, s_2) = \sum_{t=s_1}^{s_2} [x_i(t) - \tilde{\mu}_i] \mathbb{I}[I_t = i].$$

If arm i is stochastic, then $\mathbb{E}\left[\hat{D}_i(s_1, s_2)\right] = O\left(L_i(t)\tilde{\Delta}_i p_i(t)\right)$ such that a positive test suggests that the arm is non-stochastic. Since the expected number of plays of arm i is L_i^0/K in each phase,

3. The last phase ends when the total number of time steps n is exhausted or when the algorithm switches to EXP3.P.

the test is weak, though, with constant false positive and false negative rates. To avoid incorrectly classifying a stochastic arm as non-stochastic, the test is repeated several times. To make the tests independent, a new phase is started in **Step 4.b** after a detection is reported. To avoid that too much regret accumulates in the case of a non-stochastic arm, the phase length is halved. If there have been E^0 independent detections, then in **Step 4.c** there is sufficient evidence for a non-stochastic arm and the algorithm switches to EXP3.P.

In **Step 4.d** the phase ends because it has exhausted its length. Since the test in Step 4.a has given no detection, arm i has performed as expected and the algorithm has accumulated negative regret against this bad arm. This negative regret allows to start the next phase with a doubled phase length, even if the arm were non-stochastic. Doubling the phase length is necessary to avoid too many phases for a stochastic arm. (Remember that the expected number of plays of a bad arm is L_i^0/K in each phase.)

In **Step 4.e** none of the above condition is satisfied and the phase continues.

Additional simpler tests for non-stochastic arms are performed in Step 1. **Step 1.a** checks whether for all active arms the unbiased estimates of the means obey the corresponding confidence intervals. Finally, **Step 1.b** checks if the algorithm receives significantly less reward than expected from the best lower confidence bound. This may happen if a non-stochastic arm first appears close to optimal but then receives less rewards.

3.2. Choice of constants in the SAPO algorithm

In the algorithm we keep the constant names because we find them easier to read than actual values. Proper values for the constants are as follows: $C_{\mathrm{w}} = 16$, $C_{1\mathrm{b}} = 522$, $C_{\mathrm{init}} = 100/9$, $C_{\mathrm{gap}} = 60$, $C_{\mathrm{p}} = 1300$, $C_{4\mathrm{a}} = 1/10$, and $C_{\mathrm{E}} = 15$.

4. Preliminaries for the analysis of SAPO

An important tool for our analysis are concentration inequalities, in particular Bernstein's inequality for martingales and a variant of Hoeffding-Azuma's inequality for the maximum of partial sums, $\max_{1\le s\le t\le n}\sum_{i=s}^{t} Y_i$. These inequalities are given in Appendix A. We denote by $\mathcal{H}_t$ the past up to and including time t.

The next lemma states some properties of algorithm SAPO. Let

$$T_i(s_1, s_2) = \#\{t : s_1 \le t \le s_2 : I_t = i\}$$

denote the number of plays of arm i in time steps s_1 to s_2, let $n_{B,i}$ be the time when arm i is evicted from $\mathcal{A}$,

$$i \in \mathcal{A}(n_{B,i} - 1) \text{ and } i \in \mathcal{B}(n_{B,i}),$$

and let n_S be the time step when SAPO switches to EXP3.P. If SAPO never switches to EXP3.P, then $n_S = n$.

Lemma 5 (a) *If $i \in \mathcal{B}(t)$ then $\tilde{\mu}_i + \tilde{\Delta}_i < \mathrm{lcb}^*(t)$.*

(b) *For each arm the number of testing phases k, $\tau_{i,k} \cdots \tau_{i,k+1} - 1$ is at most $M = \lceil \log_2 n \rceil + 2E^0$.*

(c) *With probability $1 - O(\delta)$, the number of plays of any bad arm i is bounded as $T_i(n_{B,i}, n_S) \le \frac{101}{100} L_i^0 M/K = O\left(M/\tilde{\Delta}_i^2\right)$.*

Proof (Sketch) Statement (a) follows immediately from Step 2 of the algorithm since $\tilde{\mu}_i = \hat{\mu}_i(n_{B,i}-1)$, $\tilde{\Delta}_i = C_{\text{gap}} \cdot \text{width}_i(n_{B,i}-1)$, $\hat{\mu}_i(n_{B,i}-1) + C_{\text{gap}} \cdot \text{width}_i(n_{B,i}-1) < \text{lcb}^*(n_{B,i}-1)$, and $\text{lcb}^*(t)$ is non-decreasing.

Statement (b) follows from the fact that Step 4.b (where the phase length is halved) is executed at most E^0 times. In the other phases the phase length is doubled in Step 4.d. Since the phase length is at most n, the number of phases is at most $\log_2 n + 2E^0$.

For statement (c) we observe that by the definition of $p_i(t)$ the expected number of plays in any testing phase of a bad arm i is L_i^0/K. Thus the expected number of plays in all phases is $L_i^0 M/K$. Since the variance is bounded by the same quantity, an application of Bernstein's inequality gives the result.

Detailed proofs are given in Appendix C. ■

5. Analysis of SAPO for adversarial bandits

In this section we prove pseudo-regret bounds for SAPO against adversarial and possibly adaptive bandits. Since we know from Theorem 1 that EXP3.P suffers small regret, we only need to bound the pseudo-regret of SAPO before it switches to EXP3.P. For the remaining section we fix some arm i. We have

$$\begin{aligned}\sum_{t=1}^{n_S} x_i(t) - \sum_{t=1}^{n_S} x_{I_t}(t) &= \sum_{t=1}^{n_S} [x_i(t) - \text{lcb}^*(t)] + \sum_{t=1}^{n_S} [\text{lcb}^*(t) - x_{I_t}(t)] \\ &= \sum_{t=1}^{n_{B,i}-1} [x_i(t) - \text{lcb}^*(t)] + \sum_{t=n_{B,i}}^{n_S} [x_i(t) - \text{lcb}^*(t)] + \sum_{t=1}^{n_S} [\text{lcb}^*(t) - x_{I_t}(t)] \qquad (1)\end{aligned}$$

The first sum in (1) bounds the regret for the time when i is an active arm. For stochastic arms, the best lower confidence bound $\text{lcb}^*(t)$ would be not too far from the rewards of the arms that are still active. For non-stochastic arms, though, we need the tests in SAPO, in particular those in Step 1, to guarantee a similar behavior and achieve

$$\mathbb{E}\left[\sum_{t=1}^{n_{B,i}-1} [x_i(t) - \text{lcb}^*(t)]\right] = O\left(\sqrt{Kn\log(n/\delta)}\right), \qquad (2)$$

see Appendix D.1.

The crucial part of the analysis concerns the second sum in (1) which bounds the regret for the time when i is a bad arm. For its analysis we explicitly track negative regret to compensate for positive regret. In Section 5.1 below we sketch the main ideas for handling this sum (formal proofs are given in Appendix D.2), showing that

$$\mathbb{E}\left[\sum_{t=n_{B,i}}^{n_S} [x_i(t) - \text{lcb}^*(t)]\right] = O\left(\frac{K\log(n/\delta)}{\tilde{\Delta}_i}\right). \qquad (3)$$

Note that $1/\tilde{\Delta}_i = O\left(\text{width}_i(n_{B,i}-1)\right) = O\left(\sqrt{T_i(n_{B,i})/\log(n/\delta)}\right) = O\left(\sqrt{n/\log(n/\delta)}\right)$ such that $O\left(K\log(n/\delta)/\tilde{\Delta}_i\right) = O\left(K\sqrt{n\log(n/\delta)}\right)$.

Finally, the third sum can be observed by the algorithm and is taken care of by the test in Step 1.b, such that

$$\sum_{t=1}^{n_S} [\mathrm{lcb}^*(t) - x_{I_t}(t)] = O\left(\sqrt{Kn\log(n/\delta)}\right). \tag{4}$$

Together, inequalities (1)–(4) and the bound on EXP3.P in Theorem 1 give the bound on the pseudo-regret in Theorem 3.

5.1. Bounding the regret for bad arms

If a bad arm is non-stochastic, then it may first appear suboptimal but still be optimal after all. We need to show that the tests of our algorithm, in particular the test in Step 4.a, are sufficient to detect such a situation. Since the algorithm checks arms in $\mathcal{B}(t)$ only rarely, it will take some time for such detection. In our analysis we explicitly compensate the regret during this delayed detection by the negative regret accumulated while arm i was performing suboptimally.

We consider the testing phases k, $\tau_{i,k} \dots \tau_{i,k+1} - 1$, of arm i, and recall that L_{ik} is the length parameter for phase k and $p_{ik} = L_i^0/(KL_{ik})$ is the probability for choosing arm i in phase k. Furthermore, let E_{ik} the value of $E_i(t)$ in phase k. Note that these quantities may change only when a new phase begins. We denote by $\mathbb{P}_{ik}\{\cdot\} = \mathbb{P}\left\{\cdot|\mathcal{H}_{\tau_{i,k}-1}\right\}$ and $\mathbb{E}_{ik}[\cdot] = \mathbb{E}\left[\cdot|\mathcal{H}_{\tau_{i,k}-1}\right]$ the probabilities and expectations conditioned on the past before phase k.

For any phase we have

$$\begin{aligned}\sum_{t=\tau_{i,k}}^{\tau_{i,k+1}-1} [x_i(t) - \mathrm{lcb}^*(t)] &= \sum_{t=\tau_{i,k}}^{\tau_{i,k+1}-1} [x_i(t) - \tilde{\mu}_i + \tilde{\mu}_i - \mathrm{lcb}^*(t)] \\ &< \sum_{t=\tau_{i,k}}^{\tau_{i,k+1}-1} [x_i(t) - \tilde{\mu}_i] - \tilde{\Delta}_i[\tau_{i,k+1} - \tau_{i,k}] \end{aligned} \tag{5}$$

by Lemma 5a. Thus we want to prevent that the rewards of arm i are significantly larger than the estimated mean $\tilde{\mu}_i$. In particular, the test in Step 4.a is supposed to detect events $D_i(s_1, s_2) > 2C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik}$ with

$$D_i(s_1, s_2) := \sum_{t=s_1}^{s_2} [x_i(t) - \tilde{\mu}_i].$$

Since on average arm i is chosen only L_i^0/K times per phase, there is a constant false negative rate q_{adv} for missing such events. For appropriate C_{p}, though, the false negative rate q_{adv} is sufficiently small, $q_{\mathrm{adv}} \le 1/25$: Since $\mathbb{E}\left[\hat{D}_i(s_1, s_2)\right] = p_{ik} D_i(s_1, s_2)$ for $\tau_{i,k} \le s_1 \le s_2 < \tau_{i,k+1}$, and Step 4.a tests for $\hat{D}_i(s_1, s_2) > C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik} p_{ik}$, we can bound q_{adv} by Bernstein's inequality using that $1 \le \tilde{\Delta}_i^2 L_i^0/(KC_{\mathrm{p}})$ and a bound on the variance,

$$\mathbb{V}\left[\hat{D}_i(s_1, s_2)\right] \le L_{ik} p_{ik} = L_i^0/K \le (\tilde{\Delta}_i L_i^0/K)^2/C_{\mathrm{p}} = (\tilde{\Delta}_i L_{ik} p_{ik})^2/C_{\mathrm{p}}.$$

The formal proof is given in Lemma 13.

We use the false negative rate q_{adv} to bound $\mathbb{E}_{ik}[D_i(\tau_{i.k}, \tau_{i.k+1} - 1)]$. Each time an event $D_i(s,t) > 2C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik}$ is missed (we consider only non-overlapping such events), $D_i(\tau_{i.k}, t)$

has increased by at most $2C_{4a}\tilde{\Delta}_i L_{ik} + 1$, and the probability for the m-th miss is at most q_{adv}^m. When such an event is detected, then the phase ends and $D_i(\tau_{i.k}, t)$ again has increased by at most $2C_{4a}\tilde{\Delta}_i L_{ik} + 1$. Thus (see Lemma 15 for the formal proof)

$$\mathbb{E}_{ik}\left[D_i(\tau_{i.k}, \tau_{i.k+1} - 1)\right] \leq (2C_{4a}\tilde{\Delta}_i L_{ik} + 1) \sum_{m \geq 0} q_{\text{adv}}^m = \frac{2C_{4a}\tilde{\Delta}_i L_{ik} + 1}{1 - q_{\text{adv}}}$$

which by (5) gives

$$\mathbb{E}_{ik}\left[\sum_{t=\tau_{i,k}}^{\tau_{i,k+1}-1} [x_i(t) - \text{lcb}^*(t)]\right] < \frac{2C_{4a}\tilde{\Delta}_i L_{ik} + 1}{1 - q_{\text{adv}}} - \tilde{\Delta}_i \mathbb{E}_{ik}\left[\tau_{i,k+1} - \tau_{i,k}\right]. \tag{6}$$

Since the bound in (6) is large for large L_{ik}, we show that such a large contribution to the regret can be compensated by negative regret in previous phases due to the term $-\tilde{\Delta}_i[\tau_{i,k+1} - \tau_{i,k}]$. We show by backward induction over the phases that the expected regret starting from phase k can be bounded,

$$\mathbb{E}_{ik}\left[\sum_{t=\tau_{i,k}}^{n_S} [x_i(t) - \text{lcb}^*(t)]\right] \leq \Phi_i(k, L_{ik}) := L_{ik}\tilde{\Delta}_i/2 + 3L_i^0\tilde{\Delta}_i(M - k + 1)$$

where M is the maximal number of phases from Lemma 5b.

Lemma 6 *Let*

$$F_{ik} = \sum_{t=\tau_{i,k}}^{n_S} [x_i(t) - \text{lcb}^*(t)].$$

Then

$$\mathbb{E}_{ik}\left[F_{ik}\right] \leq \Phi_i(k, L_{ik}).$$

Proof Let k_S be the last phase before the algorithm switches to EXP3.P with $\tau_{k_S+1} - 1 = n_S$. By Lemma 5b we have $k_S \leq M$. For $k = k_S + 1$ the lemma holds trivially since $F_{i,k_S+1} = 0$.

By (6) we have

$$\mathbb{E}_{ik}\left[F_{ik}\right] \leq \frac{2C_{4a}\tilde{\Delta}_i L_{ik} + 1}{1 - q_{\text{adv}}} + \mathbb{E}_{ik}\left[F_{i,k+1} - \tilde{\Delta}_i(\tau_{i,k+1} - \tau_{i,k})\right].$$

For the expectation on the right hand side we distinguish three cases, depending on the termination condition of phase k and the value of L_{ik}.

Case 1: Phase k is terminated by the condition in Step 4.d. Then $L_{i,k+1} = 2L_{ik}$ and

$$\mathbb{E}_{ik}\left[F_{i,k+1} - \tilde{\Delta}_i(\tau_{i,k+1} - \tau_{i,k}) \middle| \textit{Case 1}\right] \leq \Phi_i(k+1, 2L_{ik}) - \tilde{\Delta}_i L_{ik} \tag{7}$$

using the induction hypothesis.

This is the case where negative regrets accumulate since $2C_{4a}/(1 - q_{\text{adv}}) < 1$.

Case 2: Phase k is terminated by the condition in Step 4.a ([4]) and $L_{ik} > L_i^0$.
Then $L_{i,k+1} = L_{ik}/2$ and

$$\mathbb{E}_{ik}\left[F_{i,k+1} - \tilde{\Delta}_i(\tau_{i,k+1} - \tau_{i,k})\middle| \textit{Case 2}\right] \le \Phi_i(k+1, L_{ik}/2). \tag{8}$$

Case 3: Phase k is terminated by the condition in Step 4.a and $L_{ik} = L_i^0$.
Then $L_{i,k+1} = L_i^0$ and

$$\mathbb{E}_{ik}\left[F_{i,k+1} - \tilde{\Delta}_i(\tau_{i,k+1} - \tau_{i,k})\middle| \textit{Case 3}\right] \le \Phi_i(k+1, L_i^0). \tag{9}$$

To complete the induction proof, we need to show that for all three cases the right hand side of (7)–(9) is upper bounded by

$$\Phi_i(k, L_{ik}) - \frac{2C_{4a}\tilde{\Delta}_i L_{ik} + 1}{1 - q_{\text{adv}}}.$$

This can be verified by straightforward calculation. ■

Now (3) follows from Lemma 6 for $k = 1$:

$$\mathbb{E}\left[\sum_{t=n_{B,i}}^{n_S} [x_i(t) - \text{lcb}^*(t)]\right] \le \Phi_i(1, L_i^0) = O\left(L_i^0 \tilde{\Delta}_i M\right) = O\left(\frac{K\log(n/\delta)}{\tilde{\Delta}_i}\right).$$

6. The stochastic analysis

In this section we assume that all arms i are indeed stochastic with means μ_i. Recall that $\Delta_i = \mu^* - \mu_i$, $\mu^* = \max_i \mu_i$. We show that with high probability the algorithm does not switch to EXP3.P and any suboptimal arm i is chosen at most $O\left(\log(n/\delta)/\Delta_i^2\right)$ times.

We already have from Lemma 5c that with probability $1 - O(\delta)$, $T_i(n_{B,i}, n_S) = O\left(M/\tilde{\Delta}_i^2\right)$ for all arms. Thus we only need to bound the number of plays before an arm is evicted from $\mathcal{A}$, $T_i(1, n_{B,i} - 1)$. The next lemma summarizes some properties of SAPO against stochastic bandits.

Lemma 7 *With probability $1 - O(\delta)$ the following holds for all time steps t and all arms i:*
(a) *If $i \in \mathcal{A}(t)$ then $|\bar{\mu}_i(t) - \mu_i| \le \overline{\text{width}}(t)/2$.*
(b) *If $i \in \mathcal{A}(t)$ then $|\hat{\mu}_i(t) - \mu_i| \le \text{width}_i(t)/2$.*
(c) *If $i \in \mathcal{A}(t)$ then $\bar{\mu}_i(t), \mu_i \in [\overline{\text{lcb}}_i(t), \overline{\text{ucb}}_i(t)]$ and $\hat{\mu}_i(t), \mu_i \ge \text{lcb}_i(t)$.*
(d) *If $\Delta_{i^*} = 0$ then $i^* \in \mathcal{A}(t)$. Furthermore, $\mu^* \ge \text{lcb}^*(t)$.*
(e) *If $i \in \mathcal{B}(t)$ then $\tilde{\Delta}_i \le 2\Delta_i$.*

Proof (Sketch) Statements (a) and (b) follow from Hoeffding-Azuma's inequality. Details are given in Appendix E.1.

For statement (c) we observe that by construction there is a time $s \le t$ with $\bar{\mu}_i(s) - \overline{\text{width}}(s) = \overline{\text{lcb}}_i(t)$. Thus (a) implies $\bar{\mu}_i(t) \ge \mu_i - \overline{\text{width}}(t)/2 \ge \bar{\mu}_i(s) - \overline{\text{width}}(s)/2 - \overline{\text{width}}(t)/2 \ge \bar{\mu}_i(s) - \overline{\text{width}}(s) = \overline{\text{lcb}}_i(t)$. The other inequalities follow analogously.

4. If k is the last phase and the phase is terminated by a condition in Step 1, then the same analysis applies but the value of $L_{k+1,i}$ is irrelevant, since $F_{i,k+1} = 0$.

Statement (d) is proven by induction on t. Let i^* be an arm with $\mu_{i^*} = \mu^*$. If $i^* \in \mathcal{A}(t-1)$ then we have by (c) that $\mu^* \geq \mathrm{lcb}^*(t-1)$. If any arm i is evicted at time t, then we have by Step 2.a and (b) that $\Delta_i = \mu^* - \mu_i \geq \mathrm{lcb}^*(t-1) - \hat{\mu}_i(t-1) - \mathrm{width}_i(t-1)/2 \geq (C_{\mathrm{gap}} - 1/2)\mathrm{width}_i(t-1) > 0$. Thus $i \neq i^*$ and $i^* \in \mathcal{A}(t)$.

This also shows that when arm i is evicted, $\tilde{\Delta}_i = C_{\mathrm{gap}} \cdot \mathrm{width}_i(t-1) \leq C_{\mathrm{gap}}/(C_{\mathrm{gap}} - 1/2)\Delta_i$, which is statement (e). ■

To get a bound on $T_i(1, n_{B,i} - 1)$, we show that $\tilde{\Delta}_i = C_{\mathrm{gap}} \cdot \mathrm{width}_i(n_{B,i} - 1)$ cannot be too small.

Lemma 8 *With probability $1 - O(\delta)$ it holds for all times t and all arms $i \in \mathcal{A}(t)$ with $T_i(t-1) \geq C_{\mathrm{init}} \log(n/\delta)$, that*

$$C_{\mathrm{gap}} \cdot \mathrm{width}_i(t-1) \geq \Delta_i/2.$$

The argument behind the lemma is that if $i \in \mathcal{A}(t)$ then $C_{\mathrm{gap}} \cdot \mathrm{width}_i(t-1) \geq \mathrm{lcb}^*(t-1) - \hat{\mu}_i(t-1)$ where $\mathrm{lcb}^*(t-1)$ is sufficiently close to μ^* and $\hat{\mu}_i(t-1)$ is sufficiently close to μ_i. The proof is given in Appendix E.2.

Since $i \in \mathcal{A}(n_{B,i} - 1)$, we get from Lemma 8 that with probability $1 - O(\delta)$,

$$T_i(n_{B,i} - 1) \leq T_i(n_{B,i} - 2) + 1 = \frac{C_{\mathrm{w}} \log(n/\delta)}{[\mathrm{width}_i(n_{B,i} - 2)]^2} + 1 \leq \frac{4 C_{\mathrm{w}} C_{\mathrm{gap}}^2 \log(n/\delta)}{\Delta_i^2} + 1.$$

Together with Lemma 5c we have with probability $1 - O(\delta)$ that for all arms,

$$T_i(n_S) \leq \frac{101}{100} L_i^0 M/K + \frac{4 C_{\mathrm{w}} C_{\mathrm{gap}}^2 \log(n/\delta)}{\Delta_i^2} + 1 = O\left(\frac{\log(n/\delta)}{\Delta_i^2}\right). \tag{10}$$

Finally, we need to bound the probability the SAPO switches to EXP3.P. Switching in Step 1.a is already handled by Lemma 7c. Switching in Step 1.b is also unlikely, since it would mean that the algorithm has accumulated large regret. This contradicts the upper bound (10). Lemma 17 shows that SAPO switches in Step 1.b only with probability $1 - O(\delta)$.

The difficult part, though, is to show that the condition in Step 4.a is not triggered too often such that Step 4.c switches to EXP3.P. We first calculate the false positive rate q_{sto}, the probability that during a given phase the condition in Step 4.a is triggered. The false positive rate is again a constant but small, $q_{\mathrm{sto}} \leq 0.21$, see Lemma 18.

Now for a fixed arm the probability that in exactly $E \geq E^0$ out of at most M phases the condition in Step 4.a is triggered, is at most $\binom{M}{E} {q_{\mathrm{sto}}}^E$. We set $p = q_{\mathrm{sto}}/(1 + q_{\mathrm{sto}})$ and use a tail bound for the binomial distribution to sum over $E = E^0, \ldots, M$:

$$\begin{aligned}\sum_{E=E^0}^{M} \binom{M}{E} {q_{\mathrm{sto}}}^E &= (1 + q_{\mathrm{sto}})^M \sum_{E=E^0}^{M} \binom{M}{E} p^E (1-p)^{M-E} \\ &\leq (1 + q_{\mathrm{sto}})^M \exp\left\{-M \cdot D(E^0/M || p)\right\}\end{aligned}$$

where $D(a||p) = a \log \frac{a}{p} + (1-a) \log \frac{1-a}{1-p}$ is the relative entropy. Since $\frac{E^0}{M} \geq \frac{C_{\mathrm{E}}}{2C_{\mathrm{E}} + 1/\log 2}$, this sum is $O(\delta/n)$ and a union bound over the arms completes the proof.

Acknowledgments

We thank the anonymous reviewers for their very valuable comments. The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 231495 (CompLACS) and from the Austrian Science Fund (FWF) under contract P 26219-N15.

References

Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2-3):235–256, 2002a.

Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic multi-armed bandit problem. *SIAM J. Comput.*, 32(1):48–77, 2002b.

Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.

Sébastien Bubeck and Aleksandrs Slivkins. The best of both worlds: Stochastic and adversarial bandits. In *COLT - The 25th Annual Conference on Learning Theory*, pages 42.1–42.23, 2012.

R. Kaas and J. M. Buhrman. Mean, median and mode in binomial distributions. *Statistica Neerlandica*, 34(1):13–18, 1980.

Shie Mannor and John N. Tsitsiklis. The sample complexity of exploration in the multi-armed bandit problem. *JMLR*, 5:623–648, 2004.

Colin McDiarmid. *Concentration*, volume 16 of *Algorithms Combin.*, pages 195–248. Springer, Berlin, 1998.

Herbert Robbins. Some aspects of the sequential design of experiments. *Bull. Amer. Math. Soc.*, 58 (5):527–535, 1952.

Yevgeny Seldin and Aleksandrs Slivkins. One practical algorithm for both stochastic and adversarial bandits. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, pages 1287–1295, 2014.

Appendix A. Concentration inequalities

Lemma 9 ((McDiarmid, 1998, Theorem 3.15)) *Let $Y_1, \ldots, Y_N$ be a martingale difference sequence with $S_N = Y_1 + \ldots + Y_N$ with the corresponding filtration $F_0 \subseteq F_1 \subseteq \ldots \subseteq F_N$. Let $Y_i \leq b$ and $\sum_{i=1}^N \mathbb{E}\left[Y_i^2 | F_{i-1}\right] \leq V$. Then for any $z \geq 0$,*

$$\mathbb{P}\left\{S_N \geq z\right\} \leq \exp\left(-z^2/(2V + 2bz/3)\right).$$

Lemma 10 ((McDiarmid, 1998, Theorem 3.13)) *Let $Y_1, \ldots, Y_N$ be a martingale difference sequence with $a_k \leq Y_k \leq b_k$ for suitable constants a_k, b_k. Then for any $z \geq 0$,*

$$\mathbb{P}\left\{\max_{1 \leq m \leq N} \sum_{k=1}^{m} Y_k \geq z\right\} \leq \exp\left(-2z^2 \Big/ \sum_{k=1}^{N} (b_k - a_k)^2\right).$$

Corollary 11 *Let $Y_1, \ldots, Y_N$ be a martingale difference sequence with $a_k \le Y_k \le b_k$ for suitable constants a_k, b_k. Then for any $z \ge 0$,*

$$\mathbb{P}\left\{\max_{1\le s\le t\le N}\sum_{k=s}^{t} Y_k \ge z\right\} \le \quad 2\exp\left(-z^2 \Big/ 2\sum\nolimits_{k=1}^{N}(b_k-a_k)^2\right).$$

Proof

$$\begin{aligned}\mathbb{P}\left\{\max_{1\le s\le t\le N}\sum_{k=s}^{t} Y_k \ge z\right\} &\le \mathbb{P}\left\{\max_{1\le t\le N}\sum_{k=1}^{t} Y_k \ge z/2\right\} + \mathbb{P}\left\{\max_{1\le s\le N}\sum_{k=1}^{s-1}(-Y_k) \ge z/2\right\}\\ &\le 2\exp\left(-z^2 \Big/ 2\sum\nolimits_{k=1}^{N}(b_k-a_k)^2\right).\end{aligned}$$

■

Appendix B. Proof of the lower bound (Theorem 2)

Let $\Delta = 1/8$. We consider a stochastic bandit problem with constant reward $x_1(t) = 1/2$ for arm 1 and Bernoulli rewards with $\mu_2 = 1/2 - \Delta$ for arm 2. We divide the time steps into phases of increasing length $L_j = 3^j\lfloor n^\alpha \rfloor$, $j = 0, \ldots, J-1$ with $J \ge \frac{1-\alpha}{\log 3}\log n$ and an incomplete last phase $j = J$. Since the pseudo-regret of the player is at most $C(\log n)^\beta$, there is a phase $j^* < J$ where the expected number of plays of arm 2 in this phase is at most B with

$$B = \frac{8C\log 3}{1-\alpha}(\log n)^{\beta-1}.$$

We construct an adversarial bandit problem by modifying the Bernoulli distribution of arm 2. Before phase j^* the distribution remains unchanged with $\mu_2 = 1/2-\Delta$, but in phase j^* and beyond we set $\mu_2 = 1/2+\Delta$. Since this bandit problem depends only on the player strategy (for identifying phase j^*) but not on the actual choices of the player, this adversary is oblivious.

Let $T_2^{j^*}$ be the number of plays of arm 2 in phase j^*, and let $\mathbb{P}_{\text{adv}}\{\cdot\}$ and $\mathbb{E}_{\text{adv}}[\cdot]$ denote the probability and expectation in respect to this adversarial bandit problem. By Lemma 12 below we have

$$\mathbb{P}_{\text{adv}}\left\{T_2^{j^*} \le 4B\right\} \ge 1/(16n^\epsilon).$$

Since $x_{I_t}(t) - \mathbb{E}_{\text{adv}}[x_{I_t}(t)|\mathcal{H}_{t-1}]$, $t = 1, \ldots, n$, forms a martingale difference sequence, we can apply Azuma-Hoeffding's inequality (Lemma 10) and obtain

$$\mathbb{P}_{\text{adv}}\left\{\sum_{t=1}^{n} x_{I_t}(t) \ge \sum_{t=1}^{n}\mathbb{E}_{\text{adv}}[x_{I_t}|\mathcal{H}_{t-1}] + \sqrt{2n\log n}\right\} \le 1/n^2$$

and

$$\mathbb{P}_{\text{adv}}\left\{T_2^{j^*} \le 4B \wedge \sum_{t=1}^{n} x_{I_t}(t) < \sum_{t=1}^{n}\mathbb{E}_{\text{adv}}[x_{I_t}|\mathcal{H}_{t-1}] + \sqrt{2n\log n}\right\} \ge 1/(16n^\epsilon) - 1/n^2. \quad (11)$$

By the construction of the adversarial bandit problem, $T_2^{j^*} \leq 4B$ implies that

$$\sum_{t=1}^{n} \mathbb{E}_{\text{adv}}\left[x_{I_t} | \mathcal{H}_{t-1}\right] \leq n/2 + 4B\Delta + (n - t_{j^*})\Delta, \tag{12}$$

where t_{j^*} denotes the time step at the end of phase j^*. For arm 2 we have

$$\begin{aligned}
\sum_{t=1}^{n} \mathbb{E}_{\text{adv}}\left[x_2(t)\right] &= n/2 - \sum_{j=0}^{j^*-1} L_j\Delta + L_{j^*}\Delta + (n - t_{j^*})\Delta \\
&= n/2 + \lfloor n^\alpha \rfloor \Delta \left(3^{j^*} - \frac{3^{j^*} - 1}{2}\right) + (n - t_{j^*})\Delta \\
&\geq n/2 + \lfloor n^\alpha \rfloor \Delta + (n - t_{j^*})\Delta.
\end{aligned}$$

Applying Azuma-Hoeffdings's inequality for arm 2 and combining with (11) and (12) we get

$$\mathbb{P}_{\text{adv}}\left\{\sum_{t=1}^{n} x_2(t) - \sum_{t=1}^{n} x_{I_t}(t) \geq \lfloor n^\alpha \rfloor \Delta - 4B\Delta - 2\sqrt{2n \log n}\right\} \geq 1/(16n^\epsilon) - 2/n^2.$$

By the condition on n, $4B\Delta \leq (\epsilon \log n)/(16\Delta)$ such that $\lfloor n^\alpha \rfloor \Delta - 4B\Delta - 2\sqrt{2n \log n} \geq n^\alpha/8 - 4\sqrt{n \log n}$, which completes the proof of the high probability lower bound.

For the lower bound on the expected regret we construct an adaptive adversary by modifying the construction above: Let $T_2^{j^*}(t)$ be the number of plays of arm 2 in phase j^* up to and including time step t. If $T_2^{j^*} = T_2^{j^*}(t_{j^*}) \leq 4B$ then the adversarial bandit problem above remains unmodified. If there is a time step $t \leq t_{j^*}$ with $T_2^{j^*}(t) > 4B$, then for all time steps $> t$ we set again $\mu_2 = 1/2 - \Delta$.

From the argument for the oblivious adversary we have

$$\begin{aligned}
&\mathbb{E}\left[\sum_{t=1}^{n} x_2(t) - \sum_{t=1}^{n} x_{I_t}(t) \middle| T_2^{j^*} \leq 4B\right] \mathbb{P}\left\{T_2^{j^*} \leq 4B\right\} \\
&\geq \left[n^\alpha/8 - 4\sqrt{n \log n}\right] \mathbb{P}\left\{T_2^{j^*} \leq 4B \wedge \sum_{t=1}^{n} x_2(t) - \sum_{t=1}^{n} x_{I_t}(t) \geq n^\alpha/8 - 4\sqrt{n \log n}\right\} \\
&\quad -n \cdot \mathbb{P}\left\{T_2^{j^*} \leq 4B \wedge \sum_{t=1}^{n} x_2(t) - \sum_{t=1}^{n} x_{I_t}(t) < n^\alpha/8 - 4\sqrt{n \log n}\right\} \\
&\geq \left[n^\alpha/8 - 4\sqrt{n \log n}\right] \left[1/(16n^\epsilon) - 2/n^2\right] - 2/n \\
&\geq \left[n^\alpha/8 - 4\sqrt{n \log n}\right] \frac{1}{(16n^\epsilon)} - 3.
\end{aligned}$$

Analogously we get

$$\begin{aligned}
&\mathbb{E}\left[\sum_{t=1}^{n} x_1(t) - \sum_{t=1}^{n} x_{I_t}(t) \middle| T_2^{j^*} > 4B\right] \mathbb{P}\left\{T_2^{j^*} > 4B\right\} \\
&\geq \quad -(4B+1)\Delta - \sqrt{2n\log n} \\
&\qquad -n \cdot \mathbb{P}\left\{T_2^{j^*} > 4B \wedge \sum_{t=1}^{n} x_1(t) - \sum_{t=1}^{n} x_{I_t}(t) < -(4B+1)\Delta - \sqrt{2n\log n}\right\} \\
&\geq \quad -(4B+1)\Delta - \sqrt{2n\log n} - 1/n \\
&\geq \quad -2\sqrt{n\log n}.
\end{aligned}$$

Thus

$$\begin{aligned}
\mathbb{E}\left[\max_i \sum_{t=1}^{n} x_i(t) - \sum_{t=1}^{n} x_{I_t}(t)\right] &\geq \left[n^{\alpha}/8 - 4\sqrt{n\log n}\right] \frac{1}{(16n^{\epsilon})} - 3 - 2\sqrt{n\log n} \\
&\geq \frac{n^{\alpha-\epsilon}}{128} - 3\sqrt{n\log n}.
\end{aligned}$$

Lemma 12 *For any n with $(\log n)^{2-\beta} \geq \frac{64C\log 3}{(1-\alpha)\epsilon}$,*

$$\mathbb{P}_{\text{adv}}\left\{T_2^{j^*} \leq 4B\right\} \geq 1/(16n^{\epsilon}).$$

Proof The proof follows a standard argument, e.g. (Mannor and Tsitsiklis, 2004).

Let $\mathbb{P}_{\text{sto}}\{\cdot\}$ and $\mathbb{E}_{\text{sto}}[\cdot]$ denote the probability and expectation in respect to the stochastic bandit problem defined above. Since $\mathbb{E}_{\text{sto}}\left[T_2^{j^*}\right] \leq B$ we have $\mathbb{P}_{\text{sto}}\left\{T_2^{j^*} > 4B\right\} < 1/4$ and thus

$$\mathbb{P}_{\text{sto}}\left\{T_2^{j^*} \leq 4B\right\} > 3/4. \tag{13}$$

Let $G_2^{j^*}$ be the sum of rewards received when playing arm 2 in phase j^*. Conditioned on $T_2^{j^*}$, $G_2^{j^*}$ is a binomial random variable with parameters $T_2^{j^*}$ and μ_2. Hence by (Kaas and Buhrman, 1980),

$$\mathbb{P}_{\text{sto}}\left\{G_2^{j^*} \leq \lfloor T_2^{j^*}(1/2-\Delta)\rfloor\right\} \leq 1/2. \tag{14}$$

Let ω denote a particular realization of rewards $x_i(t)$, $i \in \{1,2\}$, $1 \leq t \leq t_{j^*}$, and player choices $I_1, \ldots, I_{t_{j^*}}$. For any realization ω the probabilities $\mathbb{P}_{\text{sto}}\{\omega\}$ and $\mathbb{P}_{\text{adv}}\{\omega\}$ are related by

$$\begin{aligned}
\mathbb{P}_{\text{adv}}\{\omega\} &= \mathbb{P}_{\text{sto}}\{\omega\} \frac{(1/2+\Delta)^{G_2^{j^*}(\omega)}(1/2-\Delta)^{T_2^{j^*}(\omega)-G_2^{j^*}(\omega)}}{(1/2-\Delta)^{G_2^{j^*}(\omega)}(1/2+\Delta)^{T_2^{j^*}(\omega)-G_2^{j^*}(\omega)}} \\
&= \mathbb{P}_{\text{sto}}\{\omega\} \left(\frac{1-2\Delta}{1+2\Delta}\right)^{T_2^{j^*}(\omega)-2G_2^{j^*}(\omega)}.
\end{aligned}$$

If $G_2^{j^*}(\omega) \geq \lfloor (1/2-\Delta) T_2^{j^*}(\omega) \rfloor$ then

$$\begin{aligned}\mathbb{P}_{\text{adv}}\{\omega\} &\geq \mathbb{P}_{\text{sto}}\{\omega\} \left(\frac{1-2\Delta}{1+2\Delta}\right)^{T_2^{j^*}(\omega)-2((1/2-\Delta)T_2^{j^*}(\omega)-1)} \\ &= \mathbb{P}_{\text{sto}}\{\omega\} \left(\frac{1-2\Delta}{1+2\Delta}\right)^{2\Delta T_2^{j^*}(\omega)+2}.\end{aligned}$$

If furthermore $T_2^{j^*}(\omega) \leq 4B$, then

$$\mathbb{P}_{\text{adv}}\{\omega\} \geq \mathbb{P}_{\text{sto}}\{\omega\} \left(\frac{1-2\Delta}{1+2\Delta}\right)^{8\Delta B+2}.$$

Hence

$$\begin{aligned}\mathbb{P}_{\text{adv}}\left\{T_2^{j^*} \leq 4B\right\} &\geq \mathbb{P}_{\text{adv}}\left\{T_2^{j^*} \leq 4B \wedge G_2^{j^*} \geq \lfloor T_2^{j^*}(1/2-\Delta)\rfloor\right\} \\ &\geq \mathbb{P}_{\text{sto}}\left\{T_2^{j^*} \leq 4B \wedge G_2^{j^*} \geq \lfloor T_2^{j^*}(1/2-\Delta)\rfloor\right\} \left(\frac{1-2\Delta}{1+2\Delta}\right)^{8\Delta B+2} \\ &\quad \text{[by (13) and (14)]} \\ &\geq \frac{1}{4}\left(\frac{1-2\Delta}{1+2\Delta}\right)^{8\Delta B+2} \\ &\geq \frac{1}{4}(1-4\Delta)^{8\Delta B+2} \\ &\quad [\Delta = 1/8,\ 1-x \geq e^{-2x} \text{ for } 0 \leq x \leq 1/2] \\ &\geq \frac{1}{16}\exp\{-64\Delta^2 B\} \\ &\geq \frac{1}{16n^\epsilon}\end{aligned}$$

for $(\log n)^{2-\beta} \geq \frac{64C\log 3}{(1-\alpha)\epsilon}$. ■

Appendix C. Proof of Lemma 5

Proof of (b) We fix some arm i. By the condition in Step 4.c, Step 4.b can be executed for this arm at most E^0 times. Let m be the number of executions of Step 4.d for arm i, such that the number of phases is at most $m + E^0 + 1$ and the length of the longest phase is at least $2^{m-E_0-1} \cdot L_i^0$. Then $n \geq \sum_{k=1}^{m+E^0}(\tau_{i,k+1}-\tau_{i,k}) \geq 2^{m-E^0}-1+2E^0$ and $m \leq E^0+\lfloor \log_2(n-1)\rfloor \leq E^0+\lceil \log_2 n\rceil - 1$. ■

Proof of (c) We fix some arm i and use Bernstein's inequality (Lemma 9) with the martingale differences

$$Y_t = \mathbb{I}[I_t = i] - p_i(t)$$

for $n_{B,i} \le t \le n_S$ and $Y_t = 0$ otherwise. Then $Y_j \le 1$ and

$$\sum_{t=1}^{n} \mathbb{E}[Y_t^2 | \mathcal{H}_{t-1}] = \sum_{t=1}^{n} \mathbb{E}[Y_t^2 | p_i(t)] \le \sum_{t=n_{B,i}}^{n_S} p_i(t).$$

In any testing phase k, $p_i(t) = L_i^0/(KL_{ik})$ for $\tau_{i,k} \le t < \tau_{i,k+1} \le \tau_{i,k} + L_{ik}$. Thus in each phase $\sum_{t=\tau_{i,k}}^{\tau_{i,k+1}-1} p_i(t) \le L_i^0/K$ and $\sum_{t=n_{B,i}}^{n_S} p_i(t) \le L_i^0 M/K$. Hence Bernstein's inequality gives

$$\begin{aligned}\mathbb{P}\left\{T_i(n_{B,i}, n_S) \ge (1+C)L_i^0 M/K\right\} &\le \mathbb{P}\left\{\sum_{t=1}^{n} Y_t \ge CL_i^0 M/K\right\} \\ &\le \exp\left\{-\frac{C^2 L_i^0 M/K}{2 + 2C/3}\right\} \le \exp\left\{-\frac{2C^2 C_{\mathrm{p}} C_{\mathrm{E}}}{2+2C/3}\log(n/\delta)\right\} \le \delta/n\end{aligned}$$

for $C \ge 1/100$. A union bound for i completes the proof. ■

Appendix D. Proofs for SAPO against adversarial bandits

D.1. Proof of inequality (2)

We need to show that

$$\mathbb{E}\left[\sum_{t=1}^{n_{B,i}-1} [x_i(t) - \mathrm{lcb}^*(t)]\right] = O\left(\sqrt{Kn\log(n/\delta)}\right).$$

By the definition of $\bar{\mu}_i(t)$ and by Step 1.a of SAPO we have by Wald's equation that

$$\mathbb{E}\left[\sum_{t=1}^{n_{B,i}-1} x_i(t)\right] = \mathbb{E}\left[(n_{B,i}-1)\cdot\bar{\mu}_i(n_{B,i}-1)\right] \le \mathbb{E}\left[(n_{B,i}-1)\cdot\overline{\mathrm{ucb}}_i(n_{B,i}-1)\right].$$

Since $\overline{\mathrm{lcb}}_i(t) \le \mathrm{lcb}^*(t)$ and $\overline{\mathrm{ucb}}_i(t)$ is non-increasing,

$$\begin{aligned}\mathbb{E}\left[\sum_{t=1}^{n_{B,i}-1} [x_i(t) - \mathrm{lcb}^*(t)]\right] &\le \mathbb{E}\left[\sum_{t=1}^{n_{B,i}-1} \left[\overline{\mathrm{ucb}}_i(t) - \overline{\mathrm{lcb}}_i(t)\right]\right] \\ \le 2\mathbb{E}\left[\sum_{t=1}^{n_{B,i}-1} \overline{\mathrm{width}}(t)\right] &= 2\mathbb{E}\left[\sum_{t=1}^{n_{B,i}-1}\sqrt{\frac{2C_{\mathrm{w}}K\log(n/\delta)}{t}}\right] \le 4\sqrt{2C_{\mathrm{w}}Kn\log(n/\delta)}.\end{aligned}$$

D.2. Proof of inequality (3)

Lemma 13 *We fix some phase k and $s \ge \tau_{i,k}$. Let*

$$t_C(s) = \min\{s \le t < \tau_{i,k+1} : D_i(s,t) > 2C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik}\}. \tag{15}$$

If no such t exists, we set $t_C(s) = \tau_{i,k+1} - 1$. Then

$$\mathbb{P}\left\{D_i(s, t_C(s)) > 2C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik} \wedge \hat{D}_i(s, t_C(s)) < C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik} p_{ik} \,\middle|\, \mathcal{H}_{s-1}\right\} \le q_{\mathrm{adv}} := 1/25.$$

Proof We use Bernstein's inequality for martingales (Lemma 9) on the martingale differences

$$Y_t = p_{ik}[x_i(t) - \tilde{\mu}_i] - \mathbb{I}\left[I_t = i\right][x_i(t) - \tilde{\mu}_i]$$

for $s \le t \le t_C(s)$ and $Y_j = 0$ otherwise; with $b = 1$ and $V = p_{ik}L_{ik} = L_i^0/K$. We get

$$\begin{aligned}
&\mathbb{P}\left\{ D_i(s, t_C(s)) > 2C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik} \wedge \hat{D}_i(s, t_C(s)) < C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik} p_{ik} \middle| \mathcal{H}_{s-1} \right\} \\
&\le \mathbb{P}\left\{ p_{ik} D_i(s, t_C(s)) - \hat{D}_i(s, t_C(s)) > C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik} p_{ik} \middle| \mathcal{H}_{s-1} \right\} \\
&= \mathbb{P}\left\{ p_{ik} D_i(s, t_C(s)) - \hat{D}_i(s, t_C(s)) > C_{4\mathrm{a}}\tilde{\Delta}_i L_i^0/K \middle| \mathcal{H}_{s-1} \right\} \\
&\le \exp\left(-\min\{C_{\mathrm{p}}C_{4\mathrm{a}}^2/4, C_{\mathrm{p}}C_{4\mathrm{a}}/2\}\right) \le 1/25.
\end{aligned}$$

■

Lemma 14 *Consider some phase k. Then*

$$\mathbb{P}_{ik}\left\{ D_i(\tau_{i,k}, \tau_{i,k+1} - 1) \ge m(2C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik} + 1) \right\} \le q_{\mathrm{adv}}^m.$$

Proof Since $D_i(s, t+1) - D_i(s,t) \le 1$, $D_i(\tau_{i,k}, \tau_{i,k+1} - 1) \ge m(2C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik} + 1)$ implies that there are time steps $\tau_{i,k} = s_1 < s_2 < \cdots < s_{m+1} \le \tau_{i,k+1}$ with $D(s_j, s_{j+1} - 2) \le 2C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik}$ and $2C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik} < D(s_j, s_{j+1} - 1) \le 2C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik} + 1$. Furthermore, by the condition in Step 4.a, $\hat{D}_i(s_j, t) < C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik} p_{ik}$ for $j = 1, \ldots, m$ and $s_j \le t < \tau_{i,k+1}$ (otherwise the phase would have ended before $\tau_{i,k+1}$). We define the event

$$\begin{aligned}
ND_j = \{ s_{j+1} = t_C(s_j) + 1 \quad &\wedge \quad D_i(s_j, s_{j+1} - 1) > 2C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik} \\
&\wedge \quad \hat{D}_i(s_j, s_{j+1} - 1) < C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik} p_{ik} \}.
\end{aligned}$$

Then

$$\begin{aligned}
\mathbb{P}_{ik}\left\{ D_i(\tau_{i,k}, \tau_{i,k+1} - 1) \ge m(2C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik} + 1) \right\} &\le \mathbb{P}_{ik}\left\{ \bigwedge_{j=1}^{m} ND_j \right\} \\
&= \prod_{j=1}^{m} \mathbb{P}_{ik}\left\{ ND_j \middle| \bigwedge_{j'=1}^{j-1} ND_{j'} \right\} \\
&\le q_{\mathrm{adv}}^m
\end{aligned}$$

by Lemma 13. ■

Lemma 15 *For any phase k,*

$$\mathbb{E}_{ik}\left[D_i(\tau_{i,k}, \tau_{i,k+1} - 1) \right] \le \frac{2C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik} + 1}{1 - q_{\mathrm{adv}}}$$

Proof

$$\begin{aligned}
&\mathbb{E}_{ik}\left[D_i(\tau_{i,k}, \tau_{i,k+1} - 1)\right] \\
&\leq (2C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik} + 1) \sum_{m \geq 0} \mathbb{P}_{ik}\left\{D_i(\tau_{i,k}, \tau_{i,k+1} - 1) \geq m(2C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik} + 1)\right\} \\
&\leq \frac{2C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik} + 1}{1 - q_{\mathrm{adv}}}
\end{aligned}$$

by Lemma 14. ■

Appendix E. Proofs for SAPO against stochastic bandits

E.1. Proof of Lemma 7

We show that (a) and (b) hold with probability $1 - O(\delta)$. The other statements of the lemma follow from the events in (a) and (b).

Proof of (a) and (b) We fix some step t and some arm i, and condition on $T_i(t) = T$. Using Hoeffding-Azuma's inequality (Lemma 10) we find

$$\mathbb{P}\left\{\hat{\mu}_i(t) - \mu_i > \mathrm{width}_i(t)/2 | T_i(t) = T\right\} \leq \exp\left\{-C_{\mathrm{w}} \log(n/\delta)/2\right\} \leq \delta/(16Kn^2).$$

Analogously we bound $\mu_i - \hat{\mu}_i(t)$. A union bound over t, i, and T gives (b).

Since $i \in \mathcal{A}(t)$ implies $p_i(t) \geq 1/K$, Bernstein's inequality (Lemma 9) with $b = K$ and $V = Kt$ gives

$$\mathbb{P}\left\{\bar{\mu}_i(t) - \mu_i > \overline{\mathrm{width}}(t)/2\right\} \leq \exp\left\{-\frac{C_{\mathrm{w}} \log(n/\delta)}{4(2 + 2/3)}\right\} \leq \delta/(16Kn).$$

Using the same bound for $\mu_i - \bar{\mu}_i(t)$ and summing over t and i gives (a). ■

E.2. Proof of Lemma 8

Lemma 16 *With probability $1 - O(\delta)$ the following holds for all time steps t and all arms i, i': If $i' \in \mathcal{A}(t)$ and $T_i(t) \geq C_{\mathrm{init}} \log(n/\delta)$, then $T_{i'}(t) \geq T_i(t)/4$.*

Proof We fix t, i, and i'. By the construction of SAPO we have $\mathbb{P}\left\{I_t = i' | \mathcal{H}_{t-1}, i' \in \mathcal{A}(t)\right\} \geq \mathbb{P}\left\{I_t = i | \mathcal{H}_{t-1}, i' \in \mathcal{A}(t)\right\}$. From $I_s, \ldots, I_t$ we select those with $I_{t'_1}, \ldots, I_{t'_k} \in \{i, i'\}$ and define a super-martingale with differences $Y_j = \mathbb{I}\left[I_{t'_j} = i\right] - \mathbb{I}\left[I_{t'_j} = i'\right]$ for $t'_j \leq t$ and $Y_j = 0$ for $t'_j > t$.

Then

$$\begin{aligned}
&\mathbb{P}\left\{T_{i'}(t) < T_i(t)/4 \wedge i' \in \mathcal{A}(t) \wedge T_i(s,t) \geq C_{\text{init}} \log(n/\delta)\right\} \\
&= \mathbb{P}\left\{\frac{3}{8}[T_{i'}(s,t) + T_i(s,t)] < \frac{5}{8}[T_i(s,t) - T_{i'}(s,t)] \wedge i' \in \mathcal{A}(t) \wedge T_i(s,t) \geq C_{\text{init}} \log(n/\delta)\right\} \\
&\leq \sum_{k \geq C_{\text{init}} \log(n/\delta)} \mathbb{P}\left\{\frac{3}{8}k < \frac{5}{8}[T_i(s,t) - T_{i'}(s,t)] \wedge i' \in \mathcal{A}(t) \wedge T_{i'}(s,t) + T_i(s,t) = k\right\} \\
&\leq \sum_{k \geq C_{\text{init}} \log(n/\delta)} \mathbb{P}\left\{\frac{3k}{5} < \sum_{j=1}^{k} Y_j\right\} \\
&\leq \sum_{k \geq C_{\text{init}} \log(n/\delta)} \exp\left\{-\frac{9k}{50}\right\} \leq \exp\left\{-\frac{9C_{\text{init}}}{50} \log(n/\delta)\right\} \frac{1}{1 - \exp\{-9/50\}}
\end{aligned}$$

by Hoeffding-Azuma's inequality (Lemma 10). A union bound for t, i, and i' completes the proof. ■

Proof of Lemma 8
Let arm i^* be optimal, $\mu_{i^*} = \mu^*$, such that $i^* \in \mathcal{A}(t)$ by Lemma 7d. By Lemma 7b, with probability $1 - O(\delta)$ we have $|\hat{\mu}_i(t-1) - \mu_i| \leq \text{width}_i(t-1)/2$ for arms i and i^*. By construction, $\text{lcb}^*(t-1) \geq \text{lcb}_{i^*}(t-1) \geq \hat{\mu}_{i^*}(t-1) - \text{width}_{i^*}(t-1)$. By Lemma 16, with probability $1 - O(\delta)$ we have $T_{i^*}(t-1) \geq T_i(t-1)/4$. Then

$$\begin{aligned}
\Delta_i &= \mu^* - \mu_i \\
&\leq \hat{\mu}_{i^*}(t-1) + \text{width}_{i^*}(t-1)/2 - \hat{\mu}_i(t-1) + \text{width}_i(t-1)/2 \\
&\leq \text{lcb}^*(t-1) + 3\text{width}_{i^*}(t-1)/2 - \hat{\mu}_i(t-1) + \text{width}_i(t-1)/2 \\
&\leq (C_{\text{gap}} + 3 + 1/2)\text{width}_i(t-1) \\
&\leq 2C_{\text{gap}} \cdot \text{width}_i(t-1).
\end{aligned}$$

■

E.3. Considering Step 1.b

Lemma 17 *The probability that there is a time t with $\sum_{s=1}^{t-1}[\text{lcb}^*(s) - x_{I_s}(s)] > C_{1\text{b}}\sqrt{Kn\log(n/\delta)}$ is at most $O(\delta)$.*

Proof By Lemma 7d we have $\text{lcb}^*(s) \leq \mu^*$ for all s with probability $1 - O(\delta)$. Thus (10) implies that with probability $1 - O(\delta)$,

$$\begin{aligned}
&\sum_{s=1}^{t-1}(\text{lcb}^*(s) - \mathbb{E}\left[x_{I_s}(s)|\mathcal{H}_{s-1}\right]) \leq \sum_{s=1}^{t-1}(\mu^*(s) - \mathbb{E}\left[x_{I_s}(s)|\mathcal{H}_{s-1}\right]) \\
&= \sum_{i=1}^{K} \Delta_i T_i(t-1) \leq \sum_{i=1}^{K} C\frac{\log(n/\delta)}{\Delta_i} \leq \frac{C}{C_{1\text{b}}}\sqrt{Kn\log(n/\delta)}
\end{aligned}$$

for $C > \frac{101}{100}C_{\mathrm{p}}(2C_{\mathrm{E}}+1)+4C_{\mathrm{w}}C_{\mathrm{gap}}^2$. By Hoeffding-Azuma's inequality (Lemma 10) we also have

$$\mathbb{P}\left\{\max_{1\le t\le n}\sum_{s=1}^{t-1}(x_{I_s}(s)-\mathbb{E}\left[x_{I_s}(s)|\mathcal{H}_{s-1}\right])\ge\sqrt{2n\log(n/\delta)}\right\}\le(\delta/n)\le 3\delta/4.$$

Thus the lemma follows for $C_{1\mathrm{b}}\ge 522$ which satisfies $C_{1\mathrm{b}}\ge C/C_{1\mathrm{b}}+1$. ■

E.4. Considering Step 4.a

Lemma 18 *If the statements in Lemma 7 hold, then*

$$\mathbb{P}_{ik}\left\{\textit{The condition of Step 4.a is triggered for arm } i \textit{ in its phase } k\right\}\le q_{\mathrm{sto}}:=0.21.$$

Proof The probability of triggering the condition in phase k is

$$\mathbb{P}_{ik}\left\{\max_{\tau_{i,k}\le s\le t<\tau_{i,k+1}}\hat{D}(s,t)\ge C_{4\mathrm{a}}\tilde{\Delta}_i L_{ik}p_{ik}\right\}.$$

We first bound the number of plays in this round, $T_i(\tau_{i,k},\tau_{i,k+1}-1)$. Applying Bernstein's inequality (Lemma 9) with $b=1$, $V=L_{ik}p_{ik}$, and $z=L_{ik}p_{ik}$ we get

$$\mathbb{P}_{ik}\left\{T_i(\tau_{i,k},\tau_{i,k+1}-1)\ge 2L_{ik}p_{ik}\right\}\le\exp\left(-\frac{L_{ik}^2p_{ik}^2}{2L_{ik}p_{ik}+2L_{ik}p_{ik}/3}\right)\le\exp\left(-\frac{3C_{\mathrm{p}}}{8\tilde{\Delta}_i^2}\right).$$

By Lemma 7b and Step 2.b, $\mu_i-\tilde{\mu}_i\le\mathrm{width}_i(n_{B,i}-1)/2\le\tilde{\Delta}_i/(2C_{\mathrm{gap}})$. Conditioning on $T_i(\tau_{i,k},\tau_{i,k+1}-1)<2L_{ik}p_{ik}$ and applying Corollary 11 of Hoeffding-Azuma's inequality with

$$z=C_{4\mathrm{a}}\tilde{\Delta}_iL_{ik}p_{ik}-(\mu_i-\tilde{\mu}_i)2L_{ik}p_{ik}=(C_{4\mathrm{a}}-1/C_{\mathrm{gap}})\tilde{\Delta}_iL_{ik}p_{ik}$$

yields

$$\begin{aligned}&\mathbb{P}_{ik}\left\{\max_{\tau_{i,k}\le s\le t<\tau_{i,k+1}}\hat{D}(s,t)\ge C_{4\mathrm{a}}\tilde{\Delta}_iL_{ik}p_{ik}\,\middle|\,T_i(\tau_{i,k},\tau_{i,k+1}-1)<2L_{ik}p_{ik}\right\}\\&\le 2\exp\left(-\left(C_{4\mathrm{a}}-\frac{1}{C_{\mathrm{gap}}}\right)^2\frac{(\tilde{\Delta}_iL_{ik}p_{ik})^2}{4L_{ik}p_{ik}}\right)\le 2\exp\left(-\left(C_{4\mathrm{a}}-\frac{1}{C_{\mathrm{gap}}}\right)^2\frac{C_{\mathrm{p}}}{4}\right)\le 0.21\end{aligned}$$

■